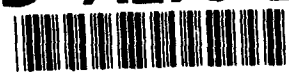

Final Report
Grant No. N00014-91-J-1102

October 1, 1991 - July 31, 1993

THE STARLITE PROJECT - PROTOTYPING REAL-TIME SOFTWARE

Submitted to:

Office of Naval Research
Chief of Naval Research
Code 1267
Ballston Tower One
Room 528
800 N. Quincy Street
Arlington, VA 22217-5660

Attention:

Dr. James G. Smith, Program Manager
Information Systems

Submitted by:

Sang H. Son
Associate Professor

SEAS Report No. UVA/525449/CS94/103
April 1994

DEPARTMENT OF COMPUTER SCIENCE

SCHOOL OF
ENGINEERING
& APPLIED SCIENCE

University of Virginia
Thornton Hall
Charlottesville, VA 22903

UNIVERSITY OF VIRGINIA
School of Engineering and Applied Science

The University of Virginia's School of Engineering and Applied Science has an undergraduate enrollment of approximately 1,500 students with a graduate enrollment of approximately 600. There are 160 faculty members, a majority of whom conduct research in addition to teaching.

Research is a vital part of the educational program and interests parallel academic specialties. These range from the classical engineering disciplines of Chemical, Civil, Electrical, and Mechanical and Aerospace to newer, more specialized fields of Applied Mechanics, Biomedical Engineering, Systems Engineering, Materials Science, Nuclear Engineering and Engineering Physics, Applied Mathematics and Computer Science. Within these disciplines there are well equipped laboratories for conducting highly specialized research. All departments offer the doctorate; Biomedical and Materials Science grant only graduate degrees. In addition, courses in the humanities are offered within the School.

The University of Virginia (which includes approximately 2,000 faculty and a total of full-time student enrollment of about 17,000), also offers professional degrees under the schools of Architecture, Law, Medicine, Nursing, Commerce, Business Administration, and Education. In addition, the College of Arts and Sciences houses departments of Mathematics, Physics, Chemistry and others relevant to the engineering research program. The School of Engineering and Applied Science is an integral part of this University community which provides opportunities for interdisciplinary work in pursuit of the basic goals of education, research, and public service.

Final Report
Grant No. N00014-91-J-1102

October 1, 1991 - July 31, 1993

THE STARLITE PROJECT - PROTOTYPING REAL-TIME SOFTWARE

Submitted to:

Office of Naval Research
Chief of Naval Research
Code 1267
Ballston Tower One
Room 528
800 N. Quincy Street
Arlington, VA 22217-5660

Attention:

Dr. James G. Smith, Program Manager
Information Systems

Submitted by:

Sang H. Son
Associate Professor

Department of Computer Science
University of Virginia
School of Engineering and Applied Science
Thornton Hall
Charlottesville, VA 22903-2442

SEAS Report No. UVA/525449/CS94/103
April 1994

Copy No._____

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	April 1994	Final Report 10/1/91 - 7/31/93

4. TITLE AND SUBTITLE
The Starlite Project - Prototyping Real-Time Software

5. FUNDING NUMBERS
N00014-91-J-1102

6. AUTHOR(S)
Sang H. Son
R. P. Cook

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Virginia
Department of Computer Science
Thornton Hall
Charlottesville, VA 22903-2442

8. PERFORMING ORGANIZATION REPORT NUMBER
UVA/525449/CS94/103

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217-5000

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The StarLite Project has the goal of discovering a set of design principles and developing efficient algorithms for distributed real-time systems. The focus of the project is on scheduling algorithms and database systems. The project also involves the construction of a prototyping environment that supports experimentation with concurrent and distributed/parallel algorithms for performance testing. One of the most important achievements in this project is the development of new scheduling algorithms based on the idea of adjusting the serialization order of active transactions dynamically. When compared with conventional transaction scheduling algorithms, our algorithms significantly improve the percentage of high priority transactions that meet the deadline. In addition, we have developed experimental database systems for performance evaluation of new technology. The RTDB server on ARTS kernel is extended to support application programmatic interface, graphic user interface, imprecise computing server, indexing mechanisms, and distributed multiple server.

14. SUBJECT TERMS	15. NUMBER OF PAGES
real-time, operating systems, database, scheduling	13
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	Unlimited

TABLE OF CONTENTS

Robert P. Cook and Sang H. Son
University of Virginia
(804) 982-2205
son@cs.virginia.edu
The StarLite Project
N00014-91-J-1102
10/1/91 - 7/31/93

1. Productivity Measures

- Refereed papers submitted but not yet published: 4
- Refereed papers published: 11
- Unrefereed reports and articles: 5
- Books or parts thereof submitted but not yet published: 3
- Books or parts thereof published: 4
- Patents filed but not yet granted: 0
- Patents granted: 0
- Invited presentations: 5
- Contributed presentations: 6
- Honors received: 11
- Prizes or awards received: 0
- Promotions obtained: 1
- Graduate students supported: 11
- Post-docs supported: 0
- Minorities supported: 0

Honors

- Son, General Chairman, Eleventh IEEE Workshop on Real-Time Operating Systems and Software, to be held in May 1994.
- Son, Program Chairman, Tenth IEEE Workshop on Real-Time Operating Systems and Software, to be held in May 1993.
- Son, Program Committee, Sixth International Conference on Distributed Computing Systems, to be held in May 1993.
- Son, Program Committee, International Symposium on Database Systems for Advanced Applications, to be held in April 1993.

- Son, ACM National Lecturer, 1991-1993.
- Son, Program Committee, Ninth IEEE Workshop on Real-Time Operating Systems and Software, 1992.
- Son, Group Leader, Ninth IEEE Workshop on Real-Time Operating Systems and Software, Group 4: Architecture, Methodology, and Databases, 1992.
- Son, Program Committee, IEEE Workshop on Transaction and Query Processing, 1992.
- Son, Program Committee, ACM SIGMOD Conference on Management of Data 1991.
- Son, Chair, Technical Activities Committee, Korean Computer Scientists and Engineers Association, 1991.
- Son, Session Chair, IEEE Real-Time Systems Symposium, 1991.
- Son, Panelist, International Conference on Very Large Data Bases (VLDB '91), on the panel "Real-Time Databases," 1991.

Promotions

Sang H. Son was promoted to the rank of Associate Professor with the granting of tenure, effective on July 1, 1992.

Graduate Students

Juhnyoung Lee (Ph.D.), scheduling real-time transactions

Shi-Chin Chiang (Ph.D.), run-time monitoring in distributed real-time databases

Young-Kuk Kim (Ph.D.), OS support for real-time database systems

Lifeng Hsu (Ph.D.), real-time distributed resource management algorithms

Henry Oh (Ph.D.), fault-tolerant multiprocessor real-time systems

Ambar Sarkar (Ph.D.), RTDB server for RS/6000 platform

Rasikan David (M.S.), real-time operating systems issues

David George (M.S.), real-time database system development

Du-Won Kang (M.S.), MRDB development

Matt Lehr (M.S.), distributed real-time testbed development

Carmen Iannacone (M.S.), RTDB development

Robert Beckinger (M.S.), temporal data modeling

Spiros Kouloumbis (M.S.), replication control in real-time databases

Stavros Yannopolous (M.S.), RTDB development

Robert P. Cook and Sang H. Son
University of Virginia
(804) 982-2205
son@cs.virginia.edu
The StarLite Project
N00014-91-J-1102
10/1/91 - 7/31/93

2. Summary of Technical Progress

Our research was directed towards discovering a set of design principles and developing efficient algorithms for distributed real-time systems and databases. Our efforts have been concentrated on three main area: real-time transaction processing, fault-tolerant multiprocessor scheduling, experimental systems and prototyping tools.

2.1. Real-Time Transactions

One of the most important achievements in this project is the development of new scheduling algorithms based on the idea of adjusting the serialization order of active transactions dynamically. This is the first successful attempt to integrate benefits of the pessimistic and optimistic approaches for transaction scheduling. Two algorithms are developed based on the notion of dynamic serialization to control blocking and aborting in a more effective manner. One is based on a priority-locking mechanism that uses the phase-dependent control of optimistic approach, while the other is based on dynamic timestamp allocation. We have implemented the first lock-based algorithm using the StarLite environment for performance evaluation. When compared with conventional transaction scheduling algorithms, it significantly improves the percentage of high priority transactions that meet the deadline. Furthermore, it is shown that the algorithm provides a very high discriminating power which enables the system to support higher priority transactions at the expense of lower priority ones when a transient overload occurs. In addition, we have evaluated optimistic concurrency control protocols for real-time database systems. Our results indicate that optimistic or hybrid approaches may outperform the pessimistic approach in a wide operational range.

We also have developed algorithms for resource management in distributed real-time systems. They are priority-ordered deadlock avoidance algorithms, efficient deadlock detection/resolution algorithms using partial resource allocation graphs, and a synchronization scheme for replicated critical data in distributed real-time database systems. Those algorithms are very efficient for distributed real-time systems, in which critical resources should be managed to support consistency, while satisfying timing constraints. Especially for replication control, we have employed a new consistency criterion, less stringent than conventional one-copy serializability. This scheme is very flexible and practical, because no prior knowledge of the data requirements or the execution time of each transaction is required. Using our StarLite prototyping environment, we have implemented those algorithms and demonstrated that they provide higher level of concurrency and greater flexibility in meeting timing requirements.

2.2. Fault-Tolerant Multiprocessor Scheduling

To investigate feasible solutions for scheduling real-time tasks in parallel/distributed environments, we have developed a new paradigm for multiprocessor real-time systems, and implemented a parallel programming interface based on our paradigm. Our new paradigm has created new research opportunities for operating systems and databases for parallel computing systems with timing and fault-tolerance requirements. For example, using the new programming interface, we have developed PRDB, an experimental real-time database system that runs on an emulated tightly-coupled multiprocessor system in the StarLite environment. It provides a general paradigm for exploiting parallelism and different real-time scheduling policies. This experimental system has been used for investigating implementation techniques for parallel database systems and the impact of multiprocessor technology on operating systems design.

To support both real-time and fault-tolerance requirements, an algorithm to schedule a number of tasks with their timing and precedence constraints on a number of processors is necessary. We have developed a scheduling model under which timing and fault-tolerance constraints can be expressed. Using this model, a scheduling problem to tolerate one arbitrary task error or processor failure has been studied. Since most multiprocessor scheduling problems are NP-complete, we have developed heuristics to obtain near-optimal solutions to the problem. We assume that all the critical tasks are periodic, and they have hard deadlines. We use two versions of each critical task, one as the primary task and the other as the secondary. The scheduling algorithm is based on the first-fit decreasing bin packing heuristics. Using the StarLite environment, the algorithm was implemented and its performance was evaluated. It was shown that the algorithm performs very well, finding the optimal solution most of the time.

2.3. Experimental Systems and Prototyping Tools

We have developed a suite of database systems on several platforms, such as StarLite, ARTS, and UNIX, and utilized them as system integration testbeds. Since a real-time system must operate in the context of operating system services, correct functioning and timing behavior of the system depends heavily on the operating system interfaces. We have developed a multi-thread database server, called RTDB, for ARTS real-time operating system kernel. The RTDB now supports application programmatic interface and graphic user interface. The application programmatic interface (API) provides an easy way for the database application programmer to construct batch clients. The API currently provides Create, Insert, Select, and Update. With imprecise server, a client can specify a deadline by which a computation (query) must complete. If the server is unable to complete the entire query, the server will return imprecise result, provided the computation had proceeded to a point where the output would be meaningful and appropriate. One problem that hinders the transformation of a non-real-time database function to a real-time one for imprecise server is recursion. Recursive function are not amenable to being stopped as easily as iterative functions. To implement the imprecise server, we have used the state machine approach in representing the execution stages of each function. Necessary actions are performed with a measurable amount of time allotted to each stage of execution.

In addition, we have developed a separate experimental system, called MRDB, a real-time database kernel running on Sun/Unix environment. In MRDB, servers and clients can be created and removed dynamically. The servers use valid time attribute and run-time estimate of requests in transaction scheduling to reduce the number of deadline-missing transactions. Using MRDB, we have performed several experiments to evaluate design alternatives in real-time scheduling and concurrency control. The temporal database kernel on Sun/Unix environment is transported to IBM RS/6000 with AIX. Ada programming interface is then developed to support a set of basic access functions to the database. We have simulated RT-DOSE (Real-time Distributed Operating System Experiments) using the interface.

Our experimental systems achieve other goal of this project—to transfer technology developed under the StarLite project to Navy, DoD, and other research organizations. Currently, Naval Ocean Systems Center in San Diego, California, is using RTDB for their distributed real-time experiments.

Robert P. Cook and Sang H. Son
University of Virginia
(804) 982-2205
son@cs.virginia.edu
The StarLite Project
N00014-91-J-1102
10/1/91 - 7/31/93

3. Publications, Presentations, and Reports

• Books and Book Chapters

(1) S. H. Son and S. Park, "Scheduling Transactions for Distributed Time-Critical Applications," in Advances in Distributed Systems, T. Casavant and M. Singhal (Editors), IEEE Computer Society, 1994(to appear).

(2) S. H. Son, J. Lee, and H. Kang, "Approaches to Design of Real-Time Database Systems," *Database Systems for Next-Genaration Applications - Principles and Practice*, W. Kim, Y. Kambayashi, and I. Paik (eds.), World Scientific Publishing, 1993.

(3) S. H. Son, C. Chang, and Y. Kim, "Performance Evaluation of Real-Time Locking Protocols," *Database Systems for Next-Genaration Applications - Principles and Practice*, W. Kim, Y. Kambayashi, and I. Paik (eds.), World Scientific Publishing, 1993.

(4) S. H. Son, R. Cook, J. Lee, and H. Oh, "New Paradigms for Real-Time Database Systems," in "Real-Time Programming," K. Ramamritham and W. Halang (Editors), Pergamon Press, 1992.

(5) R. Cook, L. Hsu, and S. H. Son, "Real-Time, Priority-Ordered, Deadlock Avoidance Algorithms," in Foundations of Real-Time Computing: Scheduling and Resource Management, A. Van Tilborg and G. M. Koob (Editors), Kluwer Academic Publishers, 1991, pp 307-324.

(6) S. H. Son, Y. Lin, and R. Cook, "Concurrency Control in Real-Time Database Systems," in Foundations of Real-Time Computing: Scheduling and Resource Management, A. Van Tilborg and G. M. Koob (Editors), Kluwer Academic Publishers, 1991, pp 185-202.

(7) R. P. Cook, "The StarLite Operating System," *Operating Systems for Mission-Critical Computing*, K. Gordon, P. Hwang, and A. Agrawala (Editors), ACM Press, 1991.

- **Refereed Journal Publications**

(8) S. H. Son, J. Ratner, S. Chiang, "StarBase: A Simulation Laboratory for Distributed Database Research," *Journal of Computer Simulation,* (to appear).

(9) S. H. Son, J. Lee, and Y. Lin, "Hybrid Protocols using Dynamic Adjustment of Serialization Order for Real-Time Concurrency Control," *Journal of Real-Time Systems,* 1992, vol. 4, no. 3, pp 269-276.

(10) S. H. Son, "Scheduling Real-Time Transactions using Priority," *Information and Software Technology,* vol. 34, no. 6, June 1992, pp 409-415.

(11) S. H. Son, "An Environment for Integrated Development and Evaluation of Real-Time Distributed Database Systems," *Journal of Systems Integration, vol. 2, no. 1, February 1992*

- **Refereed Conference Publications**

(12) S. H. Son and S. Koloumbis, "Replication Control for Distributed Real-Time Database Systems," *12th International Conference on Distributed Computing Systems,* Yokohama, Japan, pp 144-151, June 1992.

(13) S. H. Son, S. Yannopoulos, Y-K. Kim, C. Iannacone, "Integration of a Database System with Real-Time Kernel for Time Critical Applications," *Second International Conference on System Integration,* Morristown, New Jersey, pp 172-180, June 1992.

(14) S. H. Son and J. Lee, "A New Approach to Real-Time Transaction Scheduling," *4th Euromicro Workshop on Real-Time Systems,* Athens, Greece, June 1992, pp 177-182.

(15) Y. Oh and S. H. Son, "An Algorithm for Real-Time Fault-Tolerant Scheduling in Multiprocessor Systems," *4th Euromicro Workshop on Real-Time Systems* Athens, Greece, June 1992, pp 190-195.

(16) S. H. Son, S. Park, and Y. Lin, "An Integrated Real-Time Locking Protocol," *Eighth IEEE International Conference on Data Engineering,* Phoenix, Arizona, February 1992, pp 527-534.

(17) S. H. Son, J. Lee, and S. Shamsunder, "Real-Time Transaction Processing: Pessimistic, Optimistic, and Hybrid Approaches," *Second International Workshop on Transactions and Query Processing,* Tempe, Arizona, February 1992.

(18) Y. Oh and S. H. Son, "Multiprocessor Support for Real-Time Fault-Tolerant Scheduling," *IEEE Workshop on Architectural Aspects of Real-Time Systems,* San Antonio, Texas, December 1991, pp 76-80.

- **Technical Reports**

(19) Y. Oh and S. H. Son, "Fault-Tolerant Real-Time Multiprocessor Scheduling," *Technical Report TR-92-09,* Dept. of Computer Science, University of Virginia,

April 1992.

(20) S. H. Son, J. Lee, and Y. Lin, "Hybrid Protocols using Dynamic Adjustment of Serialization Order," *Technical Report TR-92-07,* Dept. of Computer Science, University of Virginia, March 1992.

(21) Y. Yang, L. Hsu, and S. H. Son, "Distributed Algorithms for Efficient Deadlock Detection and Resolution," *Technical Report TR-92-06,* Dept. of Computer Science, University of Virginia, February 1992.

• Presentations

- Son, Replication Control for Distributed Real-Time Database Systems, 12th International Conference on Distributed Computing Systems, Yokohama, Japan, June 1992.
- Son, A New Approach to Real-Time Transaction Scheduling, 4th Euromicro Workshop on Real-Time Systems, Athens, Greece, June 1992
- Son, An Integrated Real-Time Locking Protocol," IEEE International Conference on Data Engineering, Phoenix, Arizona, February 1992.
- Son, Real-Time Transaction Processing: Pessimistic, Optimistic, and Hybrid Approaches," International Workshop on Transactions and Query Processing, Tempe, Arizona, February 1992.
- Son, Multiprocessor Support for Real-Time Fault-Tolerant Scheduling, IEEE Real-Time Systems Symposium, San Antonio, Texas, December 1991.

Robert P. Cook and Sang H. Son
University of Virginia
(804) 982-2205
son@cs.virginia.edu
The StarLite Project
N00014-91-J-1102
10/1/91 - 7/31/93

4. Transitions and DOD Interactions

- Son, The Chronus Project, Proposal submitted to James Smith, Information Systems Division, Office of Naval Research, September 1992.
- Son, Replication Control for Distributed Real-Time Database Systems, presentation at the Sogang University, June 1992.
- Son, Real-Time Systems and Databases, presentation at the Seoul National University, June 1992.
- Son, the real-time database server, version 2.0, installed at NRaD distributed real-time testbed, April 1992.
- Son, real-time database project coordination meeting with IBM, Charlottesville, Virginia, April 1992.
- Son, Real-Time Systems and Databases, Kingston, Rhode Island, presentation at the University of Rhode Island, March 1992.
- Son, Real-Time Systems and Databases, presentation at the Boston University, Boston, Massachusetts, March 1992.
- Son, Advanced Real-Time Database Systems Project, Proposal submitted to Les Anderson, NRaD, February 1992.
- Son, Real-Time Database Systems, NOSC Code 413 DC^2 Review Meeting, San Diego, California, January 1992.
- Son, Real-Time Systems and Databases: Issues and Research Directions, presentation at the KSEA Symposium on Science and Technology, Washington, DC, December 1991.
- Son, StarLite project research coordination meeting with ONR, Charlottesville, Virginia, November 1991.
- Son, presentation at the ONR Foundations of Real-Time Computing Workshop, Washington, DC, October 1991.
- Son, meeting with ONR patents office staff, Charlottesville, Virginia, October 1991.
- Son, ARTS Real-Time Systems Project Review Meeting, Carnegie-Mellon University, Pittsburgh, Pennsylvania, October 1991.

Robert P. Cook and Sang H. Son
University of Virginia
(804) 982-2205
son@cs.virginia.edu
The StarLite Project
N00014-91-J-1102
10/1/91 - 7/31/93

5. Software and Hardware Prototypes

The RTDB real-time database system has been upgraded and delivered to NRaD. However, we still have a tremendous amount of work to do in fixing minor problems and identifying performance bottlenecks. The StarLite prototyping environment has been distributed to several universities as beta test sites. Both RTDB and StarLite still need a lot of work for providing proper documentation.

DISTRIBUTION LIST

1 - 3	Office of Naval Research Chief of Naval Research Code 1267 Ballston Tower One Room 528 800 N. Quincy Street Arlington, VA 22217-5660 Attention: Dr. James G. Smith, Program Manager Information Systems
4 - 9	Director Naval Research Laboratory Washington, DC 20375 Attention: Code 2627
10 - 21	Defense Technical Information Center, SR7031 Building 5, Cameron Station Alexandria, VA 22314
22	Kathy L. Raible Administrative Grants Officer Office of Naval Research Resident Representative 101 Marietta Tower, Suite 2805 101 Marietta Street Atlanta, GA 30303
23	SEAS Postaward Administration
24 - 25	S. H. Son
26	J. M. Ortega
27 - 28	H. Earnhardt, Clark Hall
29	R. A. Nixon
30	SEAS Preaward Administration Files

JO#5706:ph